ScienceLogic

Technical Brief

Fault-Tolerant Cloud Management & IT Operations with Next Generation Monitoring

TABLE OF CONTENTS

1. Introduction — About this Technical Paper

While IT management tools are crucial for building fault tolerance into today's complex and interdependent IT infrastructures, they are rarely designed or deployed with significant levels of fault tolerance themselves. This paper both discusses the critical role of IT management in data center fault tolerance and explores the unique fault tolerance capabilities designed into ScienceLogic's all-in-one solution.

2. The "Fault" with Most IT Operations Management Solutions

Have you ever thought it strange that the expression "mission-critical" is rarely used to describe systems management or IT operations management (ITOM) tools? While most IT organizations deploy a myriad of fault-tolerant technologies to ensure availability of their mission-critical applications, the tools used to manage those applications do not receive the same treatment, making them vulnerable to failure. And, typically, the fault tolerance options offered by management vendors are poorly integrated and expensive.

Some of the most common fault tolerance limitations with many ITOM products include:

Not designed with infrastructure interdependencies in mind. Fault avoidance, mitigation, detection, and resolution require monitoring and analyzing many different infrastructure components and how they interrelate. But merely being able to monitor these components is insufficient because many faults are indeterminate, requiring the correlation of numerous metrics from multiple, related components. While specialty "point" tools may report that an individual application or component is functioning within specification, a broader ITOM solution can help identify its role in a system-wide fault condition speeding up the time to resolution.

Not integrated with IT staff processes and tools. An automotive mechanic will tell you that a check engine light is useless unless the driver does something about it, and the same applies in IT operations. ITOM tools can help diagnose the root cause of faults, but resolving them will almost always involve people. Without established workflow processes and an integrated ticketing system, critical time can be wasted due to inefficient communications and lack of clarity about who is responsible. However, having these processes in place but not available can be even more devastating to the organization that depends on them in a crisis.

Expensive and hard to deploy. While many ITOM tools tout some fault tolerance capabilities, most require costly add-on modules or third-party technologies that require subject matter expertise. As a result, few ITOM tools are deployed with the fault tolerance capabilities touted on their datasheets and, since the technologies are so complex to deploy, they usually fall off the implementation plan.

Not used as part of a proactive fault tolerance strategy. Even though ITOM tools are usually key elements in disaster recovery (DR) and business continuity (BC) plans, fault tolerance for management tools is usually an afterthought. Since management tools are needed most when there is a service threat, they should have a level of fault tolerance found in a first responder toolkit — able to resist most faults, and capable of rebounding quickly from unavoidable ones.

3. ScienceLogic Fault Tolerant IT Operations Management — Instant, Intelligent, and Integrated

The ScienceLogic IT monitoring and management solution is an IT first responder's toolkit. Developed with the network operations center (NOC) in mind where fault tolerance is not an afterthought but a necessity, the solution helps identify the source of outages and reduces mean time to repair (MTTR). Unlike other ITOM tools with long deployment times and difficult learning curves, ScienceLogic is an *instant* solution that works out-of-the-box, with automation capabilities that make it *intelligent* and easy to learn and to use. Within our all-in-one solution, we also fully *integrated* all the functions and related capabilities needed for reliable, comprehensive fault tolerance, removing any need to add on packages that require additional learning curves and costs.

3.1 The ScienceLogic Philosophy

Given its service provider roots, ScienceLogic assumes that the need for advanced fault tolerant capabilities will be the norm, and not the exception. From redundant hardware, to purpose-built operating system and application components, every *integrated* aspect of a ScienceLogic system is designed to resist outages, and to recover quickly should they occur. These services have been designed to work out-of-the-box, *instantly*, without requiring deep subject matter expertise or tedious configuration, and are available by default — not subject to hidden costs or extra licenses. This *intelligent* approach — to provide "set it and forget it" fault tolerance — ensures that the product is available when needed most.

3.2 Fault-Tolerant Hardware

ScienceLogic's all-in-one solution may be packaged on physical or virtual appliances. As a physical appliance, it takes advantage of the industry's most reliable hardware that is designed to withstand many potential failures. ScienceLogic hardware includes redundant hard drives, fans, power supplies, and interfaces, all of which can be configured for fault tolerance. ScienceLogic appliances have the same level of redundancy and fault tolerance as servers hosting some of the most mission-critical applications. ScienceLogic developed and tested its software to exploit hardware fault tolerance and provide efficient and effective resiliency which has been field-proven in a variety of production environments. For further assurance, ScienceLogic appliances undergo extensive hardware and software fault tolerance testing before shipping.

3.3 Pre-Configured Appliances (Physical or Virtual)

While standardization makes sense for most IT solutions, ITOM tools don't necessarily fit in that category. Standardization can simplify the deployment and management of otherwise diverse infrastructures but can also lead to the proliferation of system threats, including operating system bugs, database vulnerabilities, viruses, firmware update problems, etc., and result in system-wide faults. Having an isolated, pre-configured physical or virtual appliance-based ITOM system ensures that it is not vulnerable to the same types of faults.

3.4 Multiple Instrumentation Methods

Standardizing on a single vendor or technology-specific method of infrastructure instrumentation is another area where standardization may

adversely impact fault tolerance. While the idea of relying upon something like VMware vCenter™ Server to abstract the complexity of the underlying infrastructure is appealing, organizations should avoid placing all of their eggs in one basket. Our recommendation is to embrace and use these new technologies while maintaining alternative means of instrumentation to avoid having one point of failure and an incomplete or skewed view of the infrastructure. Vendor-agnostic by design, ScienceLogic's solution is able to support the latest vendor-specific interfaces and to provide analysis for new technologies as well as alternative methods of instrumentation and correlation.

4. Data and Message Collection for Fault Diagnosis

Data and message collection are extremely important when diagnosing the cause of system faults. When these core components do not work, the IT staff is effectively blind to the conditions causing faults. Whether the environment is physical, virtual, or cloud-based, data collection is vital for detecting, isolating, and resolving problems. ScienceLogic's specialized expertise in comprehensive monitoring techniques allows data to be gathered using a range of methods, including passive network monitoring, native host protocols communications, proprietary third party agents, and data messages from external resources. Together, these unified collection techniques help provide a clearer picture of the actual root cause of a failure. Therefore, it is imperative that data collection is as robust and fault tolerant as possible, and fully able to withstand hardware and system failures.

4.1 Store and Retrieve Data Caching

Have you ever wondered what happens to IT monitoring data during a network outage? With most ITOM tools, that data is effectively lost forever at the very time when it is needed most. This is not the case with ScienceLogic. Each ScienceLogic data collector can operate as an autonomous resource, managing the retrieval and storage of information from remote systems and services. In the event that communication with the primary database fails, the ScienceLogic collector continues its normal operation of data collection and correlation, caching data until connectivity is reestablished. The ScienceLogic database keeps track of the last time data was synchronized for each collector.

Because of their flexibility, ScienceLogic collectors are often deployed remotely — at a network's edge, at a customer location, or in public environments. ScienceLogic collectors continue to operate despite connectivity issues, capturing critical data, both inside the network and at the point of service delivery.

4.2 Data Collection Groups

Though each individual ScienceLogic collector is fault-tolerant, there are times when data collection may be interrupted. Network outages, power failures, and rare hardware problems can result in a collector being offline.

To solve this problem, ScienceLogic's solution automatically re-provisions collector operations to other responding collectors. Referred to as "collector grouping," this capability ensures continuous operations for specific high-availability scenarios. With ScienceLogic collector groups, pairs or groups of ScienceLogic collectors can share or load-balance monitoring operations without any disruption of

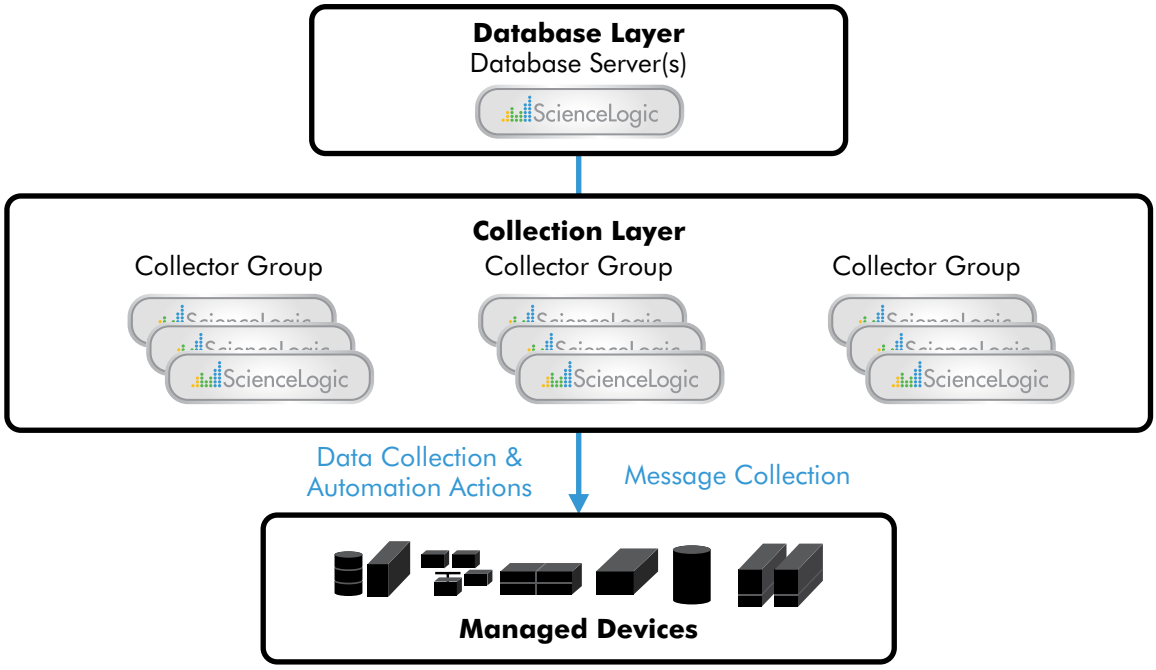

Figure 1. ScienceLogic Collector Groups ensure no loss of management data.

service or loss of data (see Figure 1). This method is highly beneficial for environments managing thousands of systems, and provides an ideal mechanism for maintenance windows or collector hardware replacement. ScienceLogic collector groups allow users to have complete administrative control over how collection fault tolerance is handled.

The flexible ScienceLogic architecture allows organizations to mix and match different collector configurations for different business needs. Each ScienceLogic solution may have a combination of multiple collector groups and numerous autonomous collectors, each configured differently. Some organizations may determine that simple mirrored collection groups are most appropriate; others may opt for a more efficient means of redundancy with a grouping of three or more collectors sharing parity across the group.

Every organization is unique and should determine which fault-tolerant collector architecture best suits its business needs. Used in conjunction with ScienceLogic's Store and Retrieve technology, ScienceLogic collector grouping provides a level of collection fault tolerance that is best-in-class.

5. Backup — the Bedrock of Fault Tolerance

Though backup technologies are ubiquitous, implementations vary wildly. The backup capabilities in ScienceLogic's platform are advanced, intuitive, and *intelligent*, and while they offer elegant failover and replication options, database backup is typically the minimum requirement. Most management products provide backup capabilities through third-party software that often is difficult and costly to implement. Some require the operational system to be brought down for a period to perform a backup

(sometimes called a "backup window"). Almost all will result in data loss when restoring from a backup. This often tolerated amount of data loss or RPO (Recovery Point Objective) is problematic because it may contain several hours or even days of mission-critical data that can hold the key to the cause of a problem. The backup capabilities in ScienceLogic are integrated in every instance, do not require a backup window, and recover data with little or no interruption to data collection.

5.1 ScienceLogic's Flexible Backup Options

ScienceLogic provides advanced database backup capabilities out-of-the-box. Easy to implement, ScienceLogic backups operate across most network-based file systems (FTP, SFTP, NFS, or CIFS/SMB), and administrators can schedule chronological backups or initiate immediate, on-demand backup jobs. ScienceLogic backups can be incremental or complete, depending on your business requirements. Further, ScienceLogic administrators can decide whether to back up all data or just mission-critical data, such as service desk information, device information, action policies, business and account information, and credentials.

5.2 ScienceLogic Backup without Downtime

ScienceLogic backups can be performed frequently with no operational disruption. These backups do not require shutting down the system, database, or collection services. With the flexibility to use any network drive for storage and without the prohibitive downtime typically associated with backups, the ScienceLogic system can automatically save local copies of configuration backups from the last seven days, and store the first configuration backup for the current and previous three months.

5.3 ScienceLogic Recovering from Backup without Data Loss

Perhaps best of all, ScienceLogic's platform can recover from a backup system with little to no loss of data due to its unique Store and Retrieve collection architecture. This "zero loss" capability allows administrators to recover from an outage while preserving the critical management data collected at the time of the outage.

6. Replication for Disaster Recovery

While backup provides a base level of fault tolerance, organizations demanding faster and more seamless recovery than even the best that backup technologies can provide will typically choose replication. ScienceLogic provides data replication out-of-the-box, *instantly*, without requiring third-party components or complex configuration. Replication provides several benefits over backup, including an extremely rapid Recovery Time Objective (RTO), little to no data loss (as described above, ScienceLogic even provides this with backup), and the ability to replicate over physically dispersed locations. Distance replication is attractive for organizations that need to safeguard against regional disasters, such as hurricanes or earthquakes.

While the benefits of replication are obvious, the cost and complexity of deploying a replicated solution is often prohibitive. Off-site disaster recovery solutions can be costly, typically leveraging existing SAN/NAS mirroring capabilities for business-critical applications. However, ScienceLogic replication capabilities are included and *integrated* with the all-in-one solution. This means that companies do not have to invest in additional SAN/NAS capabilities or limit themselves to time-

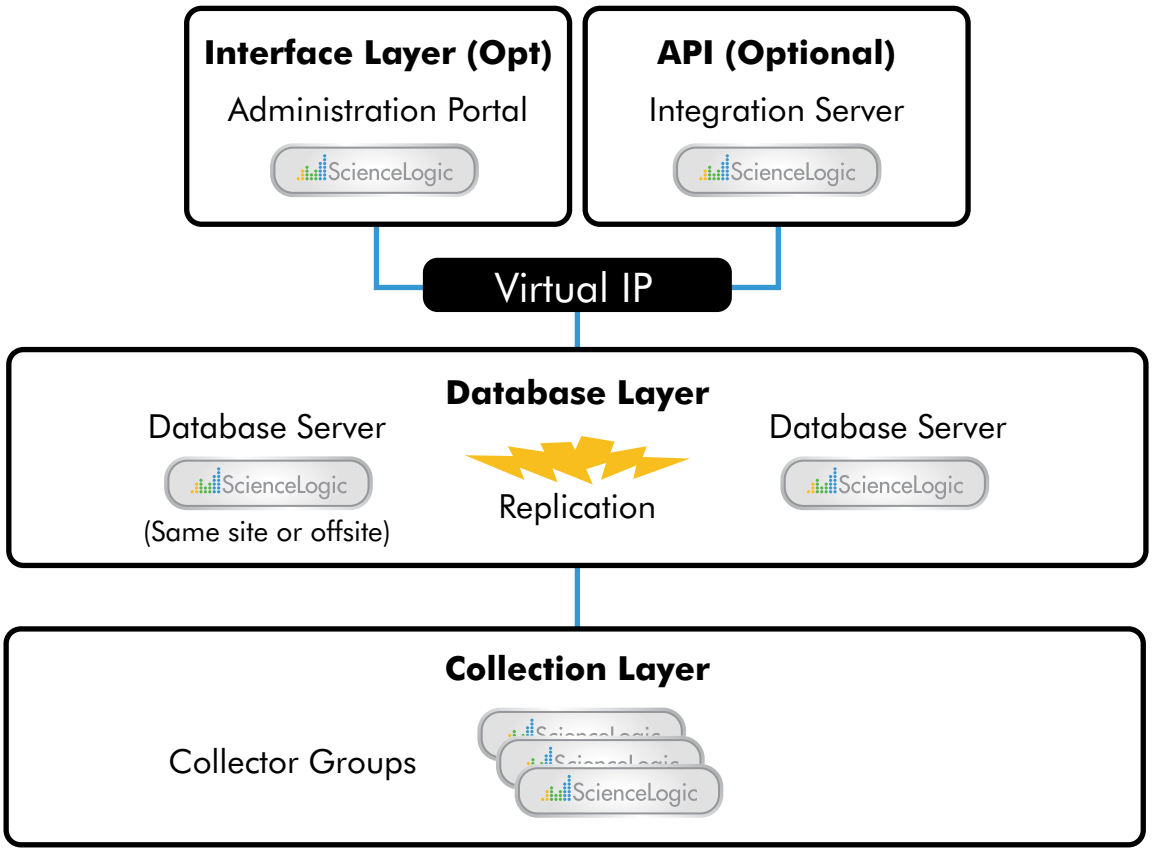

Figure 2. ScienceLogic with Database Replication.

based SnapMirrors. ScienceLogic data is replicated in real time at the device block level. Additionally, if the ScienceLogic backup database server goes off-line, replication will automatically catch up once connectivity is restored. In the event of a disaster, the ScienceLogic backup database can be enabled to take over management operations of the ScienceLogic collectors (see Figure 2).

6.1 Disaster Recovery Made Easy

Many organizations discover that replicating systems is easy, but recovering from replicated systems is hard. However, with ScienceLogic, recovering from a replicated system is easy. If the main ScienceLogic database fails, the administrator can initiate immediate take-over operations. Once executed, the recovery appliance will not only assume all incoming traffic, but will also begin retrieving data from all remote ScienceLogic collectors. In the event the take-over process was not instantaneous, data cached on the remote collectors will be retrieved to synchronize all data. In short, ScienceLogic provides continuous data resiliency despite the database being unavailable during a disaster scenario. Furthermore, when combined with Store and Retrieve collection, there is no data loss when recovering a ScienceLogic system.

6.2 Remote "Long-Distance" Replication

Even the best backup or replication solution is ineffective when all systems are impacted by the same outage. While local replication can protect organizations from a wide range of outages, they

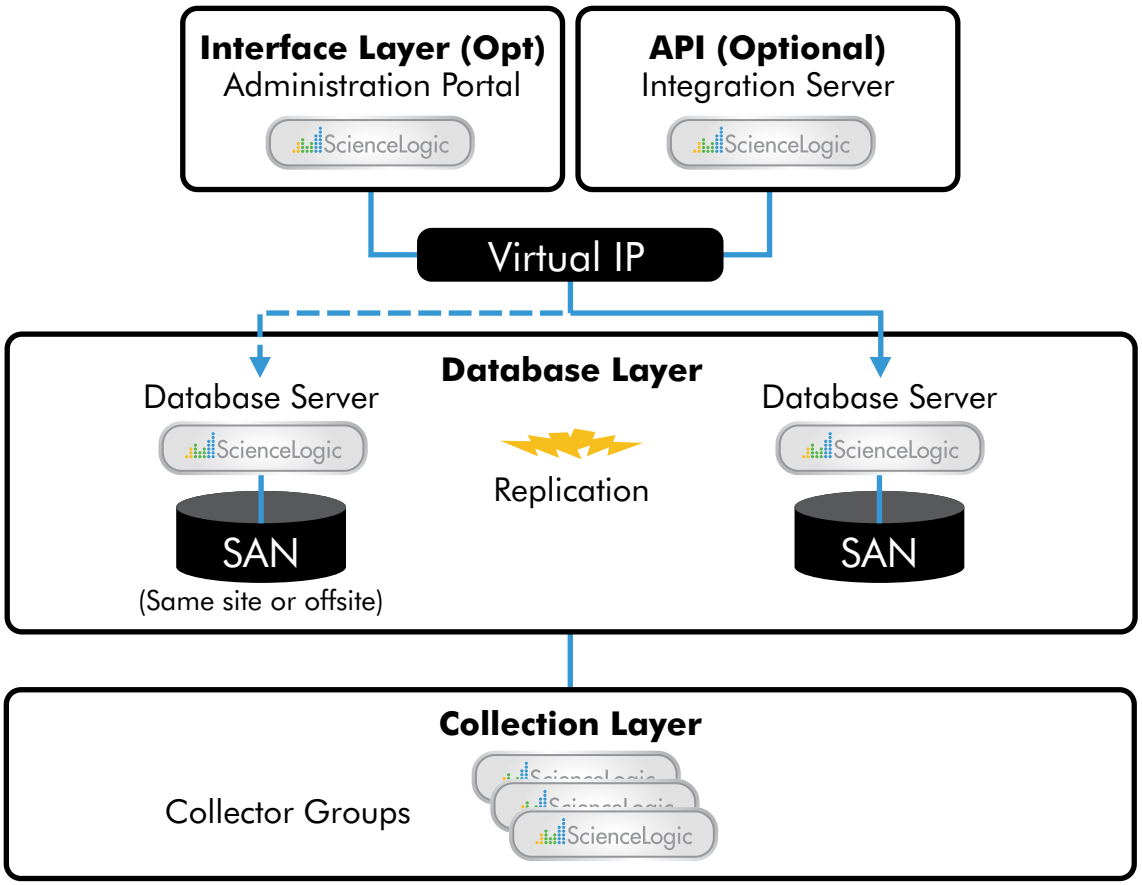

Figure 3. ScienceLogic with SAN-attached primary database and DR database.

will not provide protection against regional events (e.g., hurricanes, blackouts, earthquakes, etc.). Every ScienceLogic appliance, whether physical or virtual, ships with *integrated* support for distance replication to ensure that ScienceLogic deployments can withstand regional outages and mitigate the impact to systems, as well as resurrect them when they become available. While such long-distance replication solutions are usually prohibitively complex and expensive, every ScienceLogic system comes pre-configured with this capability. Remote replication is enabled by adding a second ScienceLogic appliance at a secure remote location. Once the ScienceLogic appliance is registered, replication begins in near-real time.

7. SAN Support

While ScienceLogic does not require networked storage for fault tolerance, many organizations depend on SAN/NAS to deliver high performance and resiliency for their mission-critical data. In support of this need, ScienceLogic database appliances provide native Fibre Channel connectivity, and support a wide range of advanced SAN-based fault tolerance technologies. As with many aspects of ScienceLogic's fault tolerance, SAN setup is simplified via a configuration wizard.

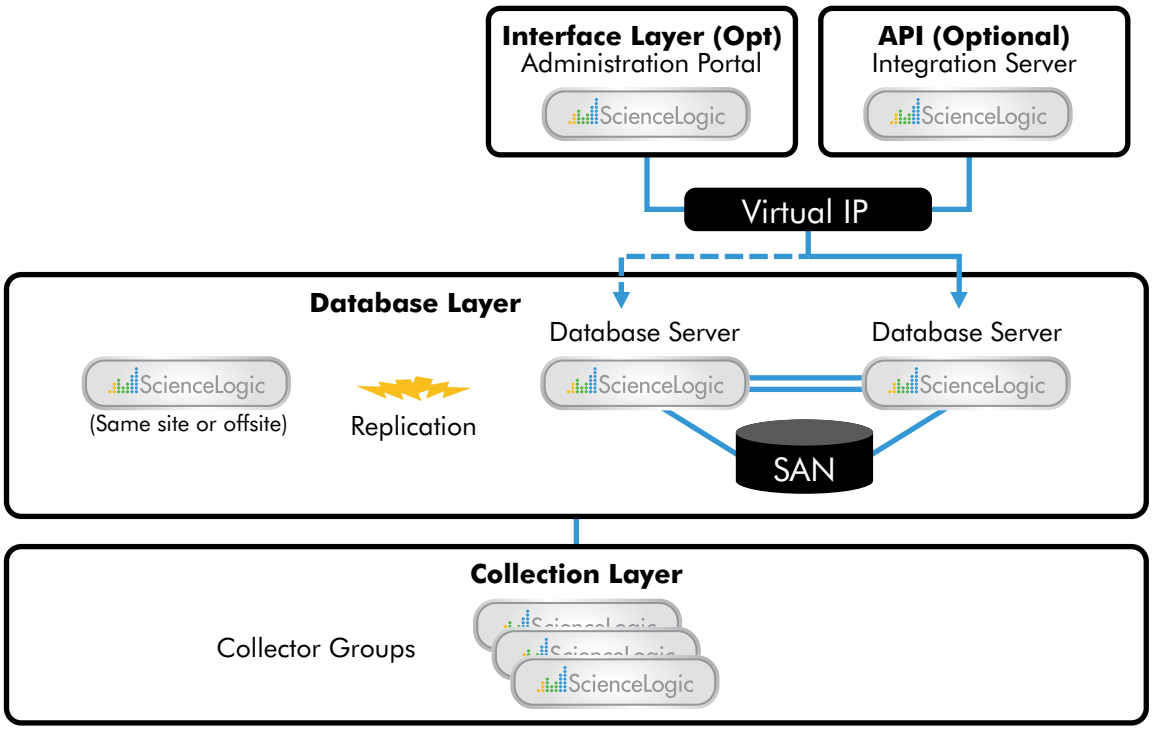

Figure 4. ScienceLogic Clustered Databases with High Availability and Disaster Recovery.

7.1 Native SAN Connectivity and Flexible SAN Services

ScienceLogic appliances can be configured with leading brand Host Bus Adapters (HBA) for native Fibre Channel communications (see Figure 3). Customers can purchase their ScienceLogic database with a pre-configured ScienceLogic validated HBA or provide their own third-party adapter to meet the needs of a specific SAN. ScienceLogic also supports several file systems for use with a third-party SAN/NAS, such as ext2, ext3, and xfs. ScienceLogic also provides large file system support for partitions above 2.1 terabytes.

7.2 SAN Multipath Support

Though SANs are reliable, even SAN connections and hardware can fail. ScienceLogic also provides native support for SAN Multipath configurations. This capability provides fault tolerance in the event of an HBA failure by enabling the mapping of two or more HBAs to a single Logical Unit Number (LUN). ScienceLogic then supports automatic failover, automatically remapping the drive connection when an individual HBA fails.

8. High-Availability Clustering

When even seconds of downtime are unacceptable, organizations deploy high availability clustering. While tedious and often costly to deploy, this is usually the technology of choice for organizations that cannot tolerate any downtime, either planned or unplanned. Most high-availability clusters are SAN-based, often requiring customization at the networking, OS, database, storage, and even application layer. Unfortunately, the complexity of properly deploying HA clustering technology often leads to failed implementations, leaving mission-critical applications exposed. Adding to this

challenge, HA clustering technology is almost never deployed for the ITOM solutions managing those applications.

8.1 ScienceLogic Makes High Availability Clustering Easy

ScienceLogic supports high-availability database clustering natively. Once the ScienceLogic system is connected to the SAN, clustering can be configured to operate in either a manual or automated failover mode (see Figure 4). The ScienceLogic configuration utility enables easy deployment of heartbeat, testing, virtual addresses, active-passive relationships, and other configuration elements that, with other tools, typically require consulting engagements or additional software.

9. Support and Field Services

We assume that every ScienceLogic system will be mission-critical, and ship every appliance with advanced customer support and logistics. For logistics, ScienceLogic partners with Dell Global Support for worldwide parts stocking and on-site field service. This relationship ensures that same-day service is available in most geographic regions and that failed components are replaced in hours, not days. Additionally, ScienceLogic support technicians are on call to assist with configuration and setup questions, defect handling, and even recovery, when necessary. This level of vendor support is a critical element that is often overlooked in fault tolerance plans.

10. ScienceLogic Delivers Flexible Fault Tolerance

Though there are a few more advanced fault tolerance options available to organizations deploying multiple, distributed ScienceLogic appliances than those deploying single-system ScienceLogic solutions, every implementation benefits from a variety of native fault tolerance HA Collection options (see Table 1). Those organizations that already have fault tolerant SANs in production will find that it is simple to integrate the ScienceLogic system with their existing SAN via an intuitive wizard. Best of all, those who start by implementing basic fault tolerance capabilities will find it easy to extend ScienceLogic's fault tolerance as their needs evolve.

	Configuration Backup	Data Backup	High Availability	Disaster Recovery	HA Collection
All-in-one	✔	✔		✔	
Distributed (Direct Disk)	✔	✔		✔	✔
Distributed (SAN)	✔	✔	✔	✔	✔

Table 1. Fault tolerance options available for various types of ScienceLogic deployments.

Like an insurance policy, fault tolerance is something that is not truly appreciated until something bad occurs. Though organizations usually have the best intentions, limited budgets and staff time often delay indefinitely the acquisition and implementation of fault tolerance technology. And, while prior outages will sometimes prompt organizations to invest in fault tolerance for their most mission-critical applications, they rarely invest in fault tolerance for the systems managing those applications. ScienceLogic, with its service provider engineering heritage, designed the its all-in-one solution to reduce this burden on IT organizations by providing easy-to-implement fault tolerance capabilities in every system.